

11-21-'05 17:00 FROM-Mattingly,Stanger 703-684-1157

T-451 P001/019 F-411

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.

ATTORNEYS AT LAW

1800 DIAGONAL ROAD, SUITE 370

ALEXANDRIA, VIRGINIA 22314

JOHN R. MATTINGLY*
DANIEL J. STANGER
SHRINATH MALUR*
CARL I. BRUNDIDGE*

COLIN D. BARNITZ
GENE W. STOCKMAN
OF COUNSEL

JEFFREY M. KETCHUM
Registered Patent Agent

* Bar Membership Other Than Virginia

PATENT, TRADEMARK
AND COPYRIGHT LAW

FACSIMILE: (703) 684-1157

RECEIVED
CENTRAL FAX CENTER

(703) 684-1120

NOV 21 2005

Date: November 21, 2005

FACSIMILE COVER LETTER

Facsimile Number: 571-273-8300

To: Special Programs Examiner Pinchus Laufer
Group Art Unit 2171, USPTO

From: Mr. Carl I. Brundidge
MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.

Re: U.S. Application No. 10/822,700
Attorney Docket No.: 500.43772X00

CERTIFICATION OF FACSIMILE TRANSMISSION

I hereby certify that the following listed documents are
being facsimile transmitted to the U.S. Patent and Trademark
Office on the date shown below:
Renewed Request for Reconsideration of Petition Under 37
C.F.R. §1.102 (MPEP §708.02) (18 pages)


Carl I. Brundidge
Reg. No. 29,621

November 21, 2005

Date

Total Number of Pages (including cover sheet): 19

If the facsimile you receive is incomplete or illegible,
please CALL (703) 684-1120. Thank you.

500.43772X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICERECEIVED
CENTRAL FAX CENTER

Applicants: N. MATSUNAMI, et al

Serial No.: 10/822,700

NOV 21 2005

Filed: April 13, 2004

For: STORAGE SYSTEM

**RENEWED REQUEST FOR RECONSIDERATION
OF PETITION UNDER 37 CFR §1.102(MPEP §708.02)****MS Petition**Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

November 21, 2005

Sir:

Applicant hereby renews its Petition to make this application Special previously submitted on August 16, 2005, in accordance with 37 CFR §1.102(d) and MPEP 708.02, VIII. The August 16, 2005 Petition was denied by a Decision issued on September 20, 2005 in which the Petitions Examiner stated that the August 16, 2005 Petition failed to recite distinct features of the claimed subject matter. The present Request for Reconsideration of Petition incorporates by reference the August 16, 2005 Petition and provides additional details regarding the claims and how the claimed subject matter is patentable over the references. The present invention is a new application filed in the United States Patent and Trademark Office on April 14, 2004 and as such has not received any examination by the Examiner.

Accordingly Applicant hereby petitions the Commissioner to make the above-identified application special in accordance with 37 CFR §1.102(d). Pursuant to MPEP §708.02(VIII), Applicant states the following.

(A) This Petition is accompanied by the fee set forth in 37 CFR §1.17(h).

The Commissioner is hereby authorized to charge any additional payment due, or to credit any overpayment, to Deposit Account No. 50-1417.

(B) All claims are directed to a single invention.

If the Office determines that all claims are not directed to a single invention, Applicant will make an election without traverse as a prerequisite to the grant of special status in conformity with established telephone restriction practice.

(C) A pre-examination search has been conducted.

The search was directed towards a storage system. In particular, the search was directed towards a first storage system being connected to a computer and a second storage system, a system for storing a file accessed by a computer including first and second storage systems and a computer and a program product for execution by the first storage system which is connected to the computer and the second storage system.

The first storage system according to the present invention includes a first storage, a first controller for controlling the first storage system and a second controller for controlling input/output operations to/from the second storage system, wherein the second storage system includes a controller and a second storage connected to the second controller and wherein the first controller creates a file system in the second storage.

The system according to the present invention includes the first and second storage system as described above wherein the first controller creates a file system in the second storage and accesses the file system created in

the second storage system in response to an access request for the file system from the computer.

The computer program product for execution by the first storage system, which is connected to the computer and the second storage system, includes a code for controlling input/output operations to/from the first storage included in the first storage system, a code for a controller included in the second storage system to control input/output operations to/from a second storage included in the second storage system, a code for creating a file system in the second storage, and a computer readable storage medium for storing the codes.

The search of the above features was conducted in the following areas:

<u>Class</u>	<u>Subclass</u>
707	200-205
709	203, 223
711	111-118, 151-154, 160-165, 170-173
714	5-7

Additionally, a computer database search was conducted on the USPTO systems EAST and WEST.

(D) The following is a list of the references deemed most closely related to the subject matter encompassed by the claims:

<u>U.S. Patent Number</u>	<u>Inventors</u>
5,537,585	Blickenstaff et al
5,659,704	Burkes et al
5,794,255	Hayashi et al
6,889,302	Gibble et al

<u>U.S. Patent Application Publication No.</u>	<u>Inventor(s)</u>
2002/0026558	Reuter et al

2002/0087751	Chong
2003/0061440	Elliott
2003/0204672	Bergsten
2004/0111580	Weber et al
2005/0071546	Delaney et al

A copy of each of these references (as well as other references uncovered during the search) is enclosed in an accompanying IDS.

(E) It is submitted that the present invention is patentable over the references for the following reasons.

It is submitted that the cited references, whether taken individually or in combination with each other, fail to teach or suggest the invention as claimed. In particular, the cited references, at a minimum, fail to teach or suggest as recited in the claims:

 a first feature of the present invention as recited in independent claim 1 of a first storage system having a first storage and a first controller, wherein the first controller creates a file system in a second storage;

 a second feature of the present invention as recited in independent claim 12 of the first controller of the first storage system creating a file system in the second storage of the second storage system and accessing the file system created in the second storage system in response to an access for the file system from the computer; and

 a third feature of the present invention as recited in independent claim 17 of a computer program for a first storage system, wherein the product includes a code for creating a file system in the second storage.

Further, the cited references fail to teach or suggest the above noted features of the present invention when taken in combination with other limitations recited in the claims.

The references considered most closely related to the claimed invention are briefly discussed below:

Blickenstaff (U.S. Patent No. 5,537,585) discloses a data storage system that is connected to a local area network and includes a storage server that on a demand basis and/or on a periodically-scheduled basis audits the activity on each volume of each data storage device that is connected to the network. Low priority data files are migrated via the network and the storage server to backend data storage media, and the directory resident in the data storage device is updated with a placeholder entry to indicate that this data file has been migrated to backend storage. When the processor requests this data file, the placeholder entry enables the storage server to recall the requested data file to the data storage device from which it originated. (See, e.g., Abstract and column 2, line 10, through column 3, line 44).

However, unlike the present invention, Blickenstaff et al. do not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Blickenstaff at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third

feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Burkes (U.S. Patent No. 5,659,704) discloses a hierarchic data storage system that stores data according to different redundancy techniques to optimize performance and reliability. The data storage system includes a disk array having a plurality of storage disks and a disk array controller which coordinates data transfer to and from the disks. The storage disks define a physical storage space. The data storage system also includes a RAID management system operatively coupled to the disk array controller for mapping two virtual storage spaces onto the physical storage space of the storage disks. The first or RAID-level virtual storage space presents the physical storage space as mirror and parity RAID areas that store data according to RAID Level 1 (mirror redundancy) and RAID Level 5 (parity redundancy). The mirror RAID areas contain multiple mirror allocation blocks and the parity RAID areas contain multiple parity allocation blocks. The second or application-level virtual storage space presents the physical storage space as multiple virtual blocks. The virtual blocks are associated with the mirror and parity allocation blocks via pointers kept in a virtual block table. (See, e.g., Abstract and column 1, line 61, through column 2, line 37).

However, unlike the present invention, Burkes does not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Burkes at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Hayashi (U.S. Patent No. 5,794,255) discloses a processing apparatus and a method of moving virtual storage resources, which moves the contents of a volume stored in accordance with the virtual storage access method to an added disk unit having a larger capacity, and to a processing apparatus and a method of moving virtual storage resources, in which movement of virtual storage resources is accomplished with a volume of the disk unit as a unit. Volume-unit transfer of information from a small-capacity disk unit to a larger-capacity disk unit is achieved to permit utilization of information including the increment capacity brought about by the transfer as virtual storage resources. The move processing comprises checking the track capacities, from among attribute information, of the source disk unit and the destination disk unit, and when these capacities are determined to be in agreement with each other, copying the volume contents of virtual storage resources of the source disk unit to the destination disk unit. After copying, the space map of the catalog section controlling the number of tracks copied to the destination is rewritten in compliance with the number of tracks of the volume capacity of the

destination. (See, e.g., Abstract and column 3, line 54, through column 4, line 18).

However, unlike the present invention, Hayashi et al. do not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Hayashi at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Gibble (U.S. Patent No. 6,889,302) discloses a method and system for maintaining information in one or more virtual volume aggregates comprising a plurality of virtual volumes. The method maintains a plurality of virtual volumes in a first information storage medium, and forms one or more virtual volume aggregates, where each of the virtual volumes is assigned to one of the virtual volume aggregates. The method further provides a plurality of second information storage media. The method then identifies the least recently used virtual volume, and writes the virtual volume aggregate comprising that least recently used virtual volume to one or more second information storage media. Thereafter, the method determines if space is required on said first information storage medium. If space is required on said

first information storage medium, the method selects the virtual volume aggregate comprising the least recently used virtual volume and removes from said first information storage medium each virtual volume in that selected virtual volume aggregate that has been written to one or more second information storage media. (See, e.g., Abstract and column 1, lines 51-63).

However, unlike the present invention, Gibble do not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Gibble at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in Independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Reuter (U.S. Patent Application Publication No. 2002/0026558) discloses a method and system for creating virtualized storage in a storage area network using distributed table-driven input/output mapping. The invention distributes the virtualization mapping in multiple parallel, mapping agents that are separate from a controller. This allows the performance-sensitive mapping process to be parallelized and distributed optimally for performance, while the control of the mapping may be located in a controller chosen for optimal cost, management, and other implementation practicalities.

The mapping agents store the virtual mapping tables in volatile memory, substantially reducing the cost and complexity of implementing the mapping agents. The controller is responsible for persistent storage of mapping tables, thereby consolidating the costs and management for persistent mapping table storage in a single component. Distributed virtualization also allows the controller to manage multiple virtual disks used by multiple host systems, and allows a single virtual disk to be shared by multiple host systems. The mapping agents preferably do not interact only with other mapping agents, thereby improving the scalability of the virtual storage system and the virtual storage system's tolerance of component failures. (See, e.g., Abstract and paragraphs 10, 21 and 22).

However, unlike the present invention, Reuter et al. do not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Reuter at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Chong (U.S. Patent Application Publication No. 2002/0087751) discloses a data storage system having a storage device controller interposed

between a host computer and one or more data storage devices where the controller manages the storage of data within the one or more storage devices. The computer system includes a data switch coupled between a host computer and one or more storage devices. The storage controller for managing the storage of data within the one or more storage devices is coupled to the switch. The switch includes a memory for storing data routing information generated by the controller, and uses the data routing information to route data directly between the host computer and the one or more storage devices such that the data does not pass through the storage controller. Within the computer system, information may be conveyed between the host computer, the switch, the one or more storage devices, and the storage controller according to a two party protocol such as the Fibre Channel protocol. The computer system achieves separation of control and data paths using a modified switch and standard host adapter hardware and host driver software. In addition, a two party protocol such as the Fibre Channel protocol is not violated. (See, e.g., Abstract and paragraph 15).

However, unlike the present invention, Chong does not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Chong at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and

further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Elliott (U.S. Patent Application Publication No. 2003/0061440) discloses hierarchical storage controllers which enable a scalable storage system. The scalable storage system is capable of causing a very large amount of discrete disk units to appear as a single drive. The scalable storage system is organized into a hierarchical structure by providing a series of disk arrays in parallel with a controller via Fibre Channel connections to form a "virtual disk". The next level in the hierarchical structure is created by connecting a number of virtual disks and a higher level controller in parallel via higher bandwidth connections. This next assemblage of devices is configured to appear as a single drive. Also, the scalable storage system disposes larger sized RAM caches to increase system performance. Moreover, wave division multiplexing (WDM) may be utilized to communicate with a file server via a highest level controller. (See, e.g., Abstract and paragraph 5).

However, unlike the present invention, Elliott does not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Elliott at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third

feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims.

Bergsten (U.S. Patent Application Publication No. 20030204672) discloses a method and system which controls data flow between one or more host processing systems and one or more data storage subsystems where the data controller provides a plurality of general and special purpose parallel processing functional units, which allows for an operator of the computer system to define the associations between the host processing systems, data controller and data storage subsystems to provide increased data protection and minimize host transfer latency. An advanced storage controller is capable of providing parallel processing capabilities to a host processing system connected storage system to increase performance, functionality and reliability of the entire computing system. The advanced storage controller comprises at least one input interface and at least one output interface, a host device simulation component, a cache device component, a physical device component and a management component. Such an advanced storage controller further, includes one or more processor elements and storage elements, which may be shared by the components or dedicated to one component. Additionally, the advanced storage controller is scalable by the static or dynamic addition of components, processors and/or memory. (See, e.g., Abstract and paragraph 16).

However, unlike the present invention, Bergsten does not disclose that said first controller copies the data in the file system created in said first

storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Bergsten at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Weber (U.S. Patent Application Publication No. 2004/0111580) discloses a method and system for managing requests of a host system to physical storage partitions. A storage system includes a plurality of storage elements with each storage element configured for providing data storage. A communications switch is communicatively connected to the storage elements for transferring requests to the physical storage partitions. A host system includes a storage router for mapping a portion of the physical storage partitions to logical storage partitions such that the host system can directly access the portion via the requests. Each of the storage elements includes a storage controller configured for processing the requests of the host system. The storage elements also include any of a disk storage device, tape storage device, CD storage device, and a computer memory storage device. The storage element controllers process received requests to access the physical storage partitions. These accesses include a variety of access types such as read and write requests to the storage partitions and control requests to

manage the storage volumes. (See, e.g., Abstract and paragraphs 14 and 31).

However, unlike the present invention, Weber et al. do not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Weber at a minimum does not teach or suggest the above described first feature of the present Invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Delaney (U.S. Patent Application Publication No. 20050071546) discloses a method and an apparatus for improving scalability of a storage system. A storage system comprising a stand-alone storage element that is reconfigurable to improve storage performance features of the storage system. The storage element comprises a plurality of disk drives, each configured for storing data. The storage element also comprises a storage controller communicatively adapted for coupling to a host computer system and configured for processing I/O requests received from the host computer system. The storage controller is also adaptable to interface with another storage controller added to the storage system. When adapted to communicate with the other storage controller, the storage controller of the

stand-alone storage element can route the I/O requests to the other storage controller through a switching fabric. In one embodiment, a first storage controller comprises: a host interface configured for communicatively coupling a host computer system to a first storage element; a storage system interface configured for communicatively coupling the first storage element to a switching fabric; and a processor configured for processing I/O requests received through the storage system interface and the host interface to access physical storage locations. The storage system interface is further configured for transferring a portion of the I/O requests through the switching fabric to a second storage controller. In another embodiment, the first storage controller is adapted to route the portion of the I/O requests to a second storage element and where the portion of the requests are processed by the second storage controller for accessing physical storage locations within the second storage element. In another embodiment, the first storage controller further comprises a disk drive interface configured for communicatively coupling to a plurality of disk drives of the first storage element to access physical storage locations of the first storage element. (See, e.g., Abstract and paragraphs 16, 17 and 27-29).

However, unlike the present invention, Delaney does not disclose that said first controller copies the data in the file system created in said first storage into said second storage, and thereafter erases the data in the file system created in said first storage.

More particularly, Delaney at a minimum does not teach or suggest the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present

invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features of the present invention in combination with the other limitations recited in each of the independent claims.

Therefore, since the cited references at a minimum fail to teach or the above described first feature of the present invention as recited in independent claim 1, the above described second feature of the present invention as recited in independent claim 12, and the above described third feature of the present invention as recited in independent claim 17, and further does not teach or suggest these features in combination with the other limitations recited in each of the independent claims, it is submitted that all of the claims are patentable over the above-described references or any of the other references of record whether taken individually or in combination with each other.

(F) Conclusion

Applicant has conducted what it believes to be a reasonable search, but makes no representation that "better" or more relevant prior art does not exist. The United States Patent and Trademark Office is urged to conduct its own complete search of the prior art, and to thoroughly examine this application in view of the prior art cited herein and any other prior art that the United States Patent and Trademark Office may locate in its own independent search. Further, while Applicant has identified in good faith certain portions of each of the references listed herein in order to provide the requisite detailed discussion of how the claimed subject matter is patentable over the

references, the United States Patent and Trademark Office should not limit its review to the identified portions but rather, is urged to review and consider the entirety of each reference, and not to rely solely on the identified portions when examining this application.

In view of the foregoing, Applicant requests that this Petition to Make Special be granted and that the application undergo the accelerated examination procedure set forth in MPEP 708.02 VIII.

(G) Fee (37 C.F.R. 1.17(i))

The fee required by 37 C.F.R. § 1.17(i) is to be paid by:

the Credit Card Payment Form (attached) for \$130.00.
 charging Account _____ the sum of \$130.00.

A duplicate of this petition is attached.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C., Deposit Account No. 50-1417 (500.43772X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.



Carl I. Brundidge
Reg. No. 29,621

CIB/jdc
(703) 684-1120